arXiv:1406.0136v2 [math.OC] 1 Jul 2014

Adaptively Blocked Particle Filtering with Spatial Smoothing in Large-Scale Dynamic Random Fields*

Francesco Bertoli † Adrian N. Bishop ‡

Abstract

The typical particle filtering approximation error is exponentially dependent on the dimension of the model. Therefore, to control this error, an enormous number of particles are required, which means a heavy computational burden that is often so great it is simply prohibitive. Rebeschini and van Handel (2013) consider particle filtering in a large-scale dynamic random field. Through a suitable localisation operation, they prove that a modified particle filtering algorithm can achieve an approximation error that is mostly independent of the problem dimension. To achieve this feat, they inadvertently introduce a systematic bias that is spatially dependent (in that the bias at one site is dependent on the location of that site). This bias consequently varies throughout field. In this work, a simple extension to the algorithm of Rebeschini and van Handel is introduced which acts to average this bias term over each site in the field through a kind of spatial smoothing. It is shown that for a certain class of random field it is possible to achieve a completely spatially uniform bound on the bias and that in any general random field the spatial inhomogeneity is significantly reduced when compared to the case in which spatial smoothing is not considered. While the focus is on spatial averaging in this work, the proposed algorithm seemingly exhibits other advantageous properties such as improved robustness and accuracy in those cases in which the underlying dynamic field is time varying.

1 Introduction

Particle filtering, as it applies here, is a powerful technique for (recursive) estimation and inference in nonlinear dynamical state-space models subject to stochastic influence. In theory, the state of an underlying stochastic dynamical system can be recursively estimated by composing the posterior probability of the state conditioned on the random observations as they become available and using a given initial prior and the stochastic dynamical system model. This process is known as recursive Bayesian filtering [1] and it is generally intractable in practice [1, 9]. The particle filter is an approximation of the Bayesian filter that employs random sampling to represent the posterior where such samples (or particles) are propagated in practice through (sequential) importance sampling [9] that attempts to capture the dynamics of the underlying system as well as the likelihood model on the observations. A resampling step inserted into the recursion is also crucial to avoid sample degeneracy and control the variance over time. To date, the particle filter has been well studied and we point to the still relevant early work [11, 12] as well as the comprehensive coverage in [9, 5] for further background.

The convergence of the particle filter in time has been well considered [8, 6]. By convergence in time, we mean that one can show that the approximation error, with respect to an idealised Bayesian filter, and due mostly to the random sampling, can be controlled through time [6]. One can even show the filter approximation error due to the sampling approximation remains bounded uniformly in time [8, 17, 7]. Such results provide significant grounding for the application of particle filtering in numerous application domains [9].

Despite substantial analysis justifying use of the particle filter in numerous applications of interest, a limitation to date surrounds application of the particle filter in high-dimensional estimation and inference problems [16, 18, 19]. The limiting factor is computational complexity. Analysis in [4, 13] suggests that the particle filter approximation error is exponential in the dimension of the underlying (measurement) model while the same error is controlled by the number of samples according to something like the inverse square root. This relationship is clearly exposited in [15]. The conclusion is simply that, an enormous number of particles (exponential in the dimension) must be maintained if one is to control the estimation error at a reasonable level when applying standard particle filter implementations in a high-dimensional dynamical estimation problem. An enormous number of particles means a heavy computational burden which is often so large it is prohibitive.

*This work was supported by AFOSR/AOARD via AOARD-144042.

†F. Bertoli is with the Australian National University and NICTA. He is supported by NICTA.

‡A.N. Bishop is with the Australian National University and NICTA. He is supported by NICTA and the Australian Research Council (ARC) via a Discovery Early Career Researcher Award (DE-120102873).

The good news is that recent studies [2, 3, 14, 15] imply that high-dimensional particle filtering may be feasible in particular applications and/or if one is willing to accept a degree of systematic bias. In [2], the particle filter is applied in a static setting where the objective is to sample from some high-dimensional target distribution. In this case, through a sequence of intermediate and simpler distributions, it is shown that the particle filter will converge to a sampled representation of the target distribution with a typical Monte Carlo error (inverse in the number of particles) given a complexity on the order of the dimension squared. Although [2] deals only, in essence, with a static problem of sampling from a fixed target distribution, the analysis introduces a novel way of thinking about high-dimensional particle filtering which may carry over to dynamic filtering problems. Related work appears in [3].

1.1 Background: The Motivating Paper

Rebeschini and van Handel in [14] consider particle filtering in large-scale dynamic random fields. They introduce a simple blocked particle filter which localises the filter to the blocks in a partition of the random field. Here, localisation means that the particle filtering prior distribution at each time is independently updated/corrected within each block through application of the observations locally conditioned on that block. The posterior over each block is independent across blocks and the posterior over the entire field is just the product of each blocked posterior. The real contribution of [14] is a descriptive and technical analysis that shows the error introduced due to the localisation procedure can be readily controlled if the dynamics of the random field at each site are only locally dependent on those sites within close proximity. The standard sampling approximation error is shown to be exponential in only the size of the individual blocks. The number of samples/particles controls the sampling approximation error at the typical rate while the error due to the localisation process is a systematic bias that can only be controlled through an increase in the block size. Since each block is updated independently, parallel implementation is readily applicable and the computational burden may be alleviated, albeit this remains to be seen in practice. While the results of [14] are at the proof-of-concept stage, the idea is incredibly powerful and it provides the entire motivation for the study herein.

It was just noted that the error due to the localisation, or blocking, procedure discussed in [14] is systematic and controlled only by the block size. Actually, this is an exaggeration and the stated result is significantly more promising. To analyse the effect of the blocking procedure, the authors introduce a tool termed the decay of correlations property which captures a spatial notion of stability, i.e. which measures how quickly dependence between sites in the field decays as a function of the distance between those sites. They show that if a suitable decay of correlations exists, then the error at any site in the field introduced due to the blocking procedure alone is dependent primarily on that site's distance to the border of the block. In other words, the systematic bias introduced due to the blocking operator is small at those sites in the field far removed from the block borders, large on the borders and varies in between. The blocking bias is not spatially uniform. The error due to the sampling procedure inherent to all particle filters is still exponential in the block size and controlled by the number of particles. The sampling approximation error implies one should not seek excessively large block sizes. On the other hand, small blocks imply the spatial inhomogeneity of the total error caused by the blocking bias is exasperated. Ideally, one would like an algorithm that maintains a spatially homogeneous error within each block. One can then focus on considering, in detail, the block-size vs. error tradeoff. The authors in [14] devote much discussion to the spatial inhomogeneity of the blocking bias and its significance.

1.2 Contribution

This work details a simple, yet relevant, algorithmic adjustment to the blocked particle filter introduced in [14]. The idea is to spatially average the bias caused by the blocking procedure by considering an adaptive sequence of partitions over the random field instead of a single partition. This adaptive blocking procedure has the effect of spatially smoothing the error caused by any single application of the blocking/localisation operation. In a certain class of random fields, this smoothing effect leads to a completely spatially homogeneous error bound. That is, the bound on the total filtering error at any site in the field is completely independent of the location of that site. In more general random fields, this spatial averaging effect can lead to a significant reduction in the spatial inhomogeneity of the particle filter error bounds when compared with [14]. This smoothing effect is of practical relevance in those cases in which the blocks should be kept relatively small for computational reasons or in which the dynamics of the random field are less than trivially localised.

The results presented here are largely at the same proof-of-concept level as those presented in the motivating paper [14]. The stated results in both studies are of a quantitative nature that is far from optimal. Nevertheless, the applicability of the particular algorithms is likely far less restricted than a strict reading of the results would suggest; indeed this is seemingly also true for the celebrated time-uniform particle filter con-

vergence results [8]. It is with this applicability in mind that the algorithmic extensions considered herein are proposed. The extensions considered are conceptually simple, easy to implement, and do not generally add to the computational burden of the algorithm. Beyond the important spatial smoothing effect of the proposed filter, allowing for (adapting) multiple partitions of the field may also provide algorithmic robustness in those cases in which the random field is time-varying etc as the partitions can be adapted online, or it may allow one to adaptively focus computation on certain locations of interest for periods of time etc. Other advantages of partition adaptation are envisioned.

1.3 Paper Organization

The paper is organised as follows. In Section 2 we introduce the model and problem setup. In Section 3 we introduce the Bayesian filtering framework, the (standard) particle filtering algorithm and we introduce the adaptively blocked particle filter for high-dimensional estimation problems. In Section 4 we state the main result, which consists of a time-average and spatially smoothed total error bound on the adaptively blocked particle filter approximation to the true Bayesian nonlinear filter. In Section 4 we also note that the bound may well be completely spatially-uniform and we outline the strategy for proving this result. In Section 5 we explore in more detail the error bound and the spatial smoothing effect of adaptively sequencing through partitions in the blocked particle filter. In Section 5 we discuss, in more detail, the class of random fields and the sequence of partitions that may lead to complete spatial uniformity in the error bound.

To this point, the problem formulation, algorithm, convergence results, and the algorithmic/convergence discussions are given. A casual reader may stop at this point, and take for granted the convergence results and the (increased) spatial homogeneity of the filtering approximation error across the random field. Indeed, the conceptual simplicity of the algorithm, and its spatial averaging property, may be sufficient to convince one of the general existence of such a convergence result (given also the results and analysis in [14]). Subsequently, the details of this result, as they are far from optimal, may be of lesser significance.

Going forward, in Section 6 we provide the technical analysis leading to the main result and following the proof strategy introduced earlier. The proofs required in this work largely overlap with those in [14] and only the required changes are derived here, with reference made to the motivating paper as often as possible. Indeed, we encourage all readers interested in high-dimensional particle filtering to study [14] since a very descriptive and accessible coverage of this topic and the blocked particle filter is provided therein (prior to the detailed technical analysis set out in [14] and to which we point as often as possible to prove our case).

2 Problem Setup

For simplicity and to ease comparison we borrow the problem scenario directly from [14].

Consider a Polish state space $\mathbb{X}$ with σ-algebra $\mathscr{X}$ and reference measure ψ, and a Polish state space $\mathbb{Y}$ with σ-algebra $\mathscr{Y}$ and a reference measure φ. Introduce on $\mathbb{X}$, a Markov chain $(X_n)_{n\geq 0}$ with a transition density $p : \mathbb{X} \times \mathbb{X} \to \mathbb{R}_+$ with respect to ψ. Introduce on $\mathbb{Y}$, a sequence $(Y_n)_{n\geq 0}$ that is conditionally independent given $(X_n)_{n\geq 0}$ and has a transition density $g : \mathbb{X} \times \mathbb{Y} \to \mathbb{R}_+$ with respect to φ. We interpret $(X_n)_{n\geq 0}$ as an underlying dynamical process that is observed through $(Y_n)_{n\geq 0}$. The pair $(X_n, Y_n)_{n\geq 0}$ is also a Markov chain.

Now suppose the state (X_n, Y_n) at each time n is a random field $(X_n^v, Y_n^v)_{v\in V}$ indexed by a (finite) undirected graph $G = (V, E)$ where V corresponds to the set of sites and E corresponds to the set of edges that define the structure of the field. The dimension of the model is then at least as big as the cardinality of the vertex set V and we assume this to be large.

To be more precise, the space $\mathbb{X}$ and $\mathbb{Y}$ are of product form $\mathbb{X} = \prod_{v\in V} \mathbb{X}^v$ and $\mathbb{Y} = \prod_{v\in V} \mathbb{Y}^v$ respectively. The associated reference measures then given by $\psi = \bigotimes_{v\in V} \psi^v$ and $\varphi = \bigotimes_{v\in V} \varphi^v$ where ψ^v and φ^v are reference measures on $\mathbb{X}^v$ and $\mathbb{Y}^v$ respectively. The transition densities p and g are given by

$$p(x, z) = \prod_{v\in V} p^v(x, z^v), \quad g(x, y) = \prod_{v\in V} g^v(x^v, y^v)$$

where $p^v : \mathbb{X} \times \mathbb{X}^v \to \mathbb{R}_+$ and $g^v : \mathbb{X}^v \times \mathbb{Y}^v \to \mathbb{R}_+$ are defined with respect to ψ^v and φ^v respectively.

The model assumes the observations $(Y_n)_{n\geq 0}$ are completely local in the sense that $g^v(x^v, y^v)$ depends only on x^v or, in other words, the conditional distribution of Y_n^v given X_n depends only on X_n^v.

A distance $d(v, v')$, that counts hops along the shortest path between $v, v' \in V$, is associated with $G = (V, E)$. Now for a fixed $r \in \mathbb{N}$ and for each vertex $v \in V$ we define $N(v) = \{v' \in V : d(v, v') \leq r\}$ which specifies a neighbourhood of v. We then assume the dynamics of $(X_n)_{n\geq 0}$ are local in the sense that $p^v(x, z^v)$ depends only on $x^{N(v)}$ or in other words the conditional distribution of X_n^v given $X_0, \ldots, X_{n-1}$ depends only on $X_{n-1}^{N(v)}$. More precisely, the dynamics obey $p^v(x, z^v) = p^v(\tilde{x}, z^v)$ whenever $x^{N(v)} = \tilde{x}^{N(v)}$ where $x^J = (x^j)_{j\in J}$ for $J \subseteq V$.

We refer to the motivating paper [14] for further discussion on such models and the references therein for background of where such models appear in the literature. Also, see [20] for such modelling motivation.

Since the process $(X_n)_{n\geq 0}$ is not directly observable, the filtering problem of interest is one of recursively estimating the unobserved state X_n given the observation history $Y_1,\ldots,Y_n$. That is, the filtering problem is one of computing

$$\pi_n^\mu \triangleq \mathbf{P}^\mu[X_n \in \cdot\,|Y_1,\ldots,Y_n]$$

where $\mathbf{P}^\mu$ is the probability measure under which $(X_n,Y_n)_{n\geq 0}$ is a Markov chain with a transition probability P that can be factored as $P((x,y),A) = \int \mathbf{1}_A(x',y')p(x,x')g(x',y')\psi(dx')\varphi(dy')$ for any $A \in \mathscr{X}\times\mathscr{Y}$ and where the initial condition $X_0 \sim \mu$ is an arbitrary probability measure μ on $\mathbb{X}$. Notationally, $\pi_n^x \triangleq \pi_n^{\delta(x)}$.

3 Adaptively Blocked Particle Filtering

Firstly, we outline the ideal nonlinear recursive Bayes filter, followed by the standard bootstrap particle filter. We note briefly the computational problem involved in applying the bootstrap filter to high-dimensional estimation problems. We then outline our adaptively blocked particle filter and note its straightforward relationship to the algorithm of [14] and discuss generally the motivation for this algorithm as it applies to filtering of high-dimensional systems.

It is well known that through an application of Bayes rule, the filter π_n^μ can be computed recursively via

$$\pi_0^\mu = \mu, \quad \pi_n^\mu = \mathsf{F}_n\pi_{n-1}^\mu \quad (n\geq 1)$$

where, it is common for practical, as well as conceptual, reasons to define $\mathsf{F}_n \triangleq \mathsf{C}_n\mathsf{P}$ where

$$(\mathsf{P}\rho)(f) \triangleq \int f(x')\,p(x,x')\,\psi(dx')\,\rho(dx)$$

is a prediction in which the filter estimate π_{n-1}^μ is propagated forward using the dynamics of the underlying process $(X_n)_{n\geq 0}$ and

$$(\mathsf{C}_n\rho)(f) \triangleq \frac{\int f(x)\,g(x,Y_n)\,\rho(dx)}{\int g(x,Y_n)\,\rho(dx)}$$

is a so-called correction (or update) in which the predicted distribution is updated by conditioning it on the observation Y_n to obtain π_n^μ. Graphically,

$$\pi_{n-1}^\mu \xrightarrow{\text{prediction}} \pi_{n-}^\mu = \mathsf{P}\pi_{n-1}^\mu \xrightarrow{\text{correction}} \pi_n^\mu = \mathsf{C}_n\pi_{n-}^\mu$$

The recursive structure of the filter allows estimation of the underlying dynamic process to be carried out 'online' over a long time horizon and incorporating measurements as they become available. However, it is well known that to this point such a filter is impractical since at the level of arbitrary probability measures it must be considered of infinite dimension. In general, no exact finite dimensional nonlinear filter can be computed.

One approximation to the nonlinear filter employs sampling and Monte Carlo approximation. Let $N\geq 1$ denote the number of samples (or particles) used in the approximation and define S^N to be the sampling operator which computes a random measure

$$\mathsf{S}^N\rho \triangleq \frac{1}{N}\sum_{i=1}^N \delta_{x(i)}, \quad x(i) \text{ is i.i.d.} \sim \rho$$

with respect to some probability measure ρ. This random measure is a discrete approximation of ρ and converges to ρ with N at a typical rate of $1/\sqrt{N}$.

The most common and arguably the simplest Monte Carlo approximation of nonlinear filtering is given by

$$\hat{\pi}_0^\mu = \mu, \quad \hat{\pi}_n^\mu = \hat{\mathsf{F}}_n\hat{\pi}_{n-1}^\mu \quad (n\geq 1)$$

where $\hat{\mathsf{F}}_n \triangleq \mathsf{C}_n\mathsf{S}^N\mathsf{P}$ now consists of three operations

$$\hat{\pi}_{n-1}^\mu \xrightarrow{\text{prediction}} \mathsf{P}\hat{\pi}_{n-1}^\mu \xrightarrow{\text{sampling}} \hat{\pi}_{n-}^\mu = \mathsf{S}^N\mathsf{P}\hat{\pi}_{n-1}^\mu \xrightarrow{\text{correction}} \hat{\pi}_n^\mu = \mathsf{C}_n\hat{\pi}_{n-}^\mu$$

This recursion yields the bootstrap particle filtering algorithm [11]. This algorithm is simple to implement and $\hat{\pi}_n^\mu$ converges to the exact filter π_n^μ as $N\to\infty$. Resampling and other operations are typically incorporated into the bootstrap particle filter to improve performance [9].

If the (standard) bootstrap particle filter is applied to a system of dimension $|V|$, then, typically [16, 4, 13], the approximation error is exponential in $|V|$ and inversely proportional to something like $\sqrt{N}$. If $|V|$ is large, then one needs a huge number of particles N to achieve a desired error rate and this requires a heavy computational burden. In many applications, like target tracking [9], there are typically no computational barriers to achieving an acceptable error. In large-scale, high-dimensional, estimation problems the particle filter is often computationally infeasible which motivates the study in [14] and obviously in this work.

To this end, we introduce a partition $\mathscr{K}$ of the vertex set V into non-overlapping blocks

$$K \cap K' = \emptyset,\ K \neq K',\ K, K' \in \mathscr{K} \text{ and } V = \bigcup_{K \in \mathscr{K}} K$$

Now suppose there exists a finite number $m \in \mathbb{N}$ of partitions $\mathscr{K}_0, \ldots, \mathscr{K}_{m-1}$ of this type. There exists a non-negative constant $\theta \in \mathbb{R}_+$ and a positive $\vartheta \in \mathbb{R}_+$ such that given a positive $\beta \in \mathbb{R}_+$ then for every node $v \in V$ we have

$$\theta \leq \theta_m(v) = \frac{1}{m} \sum_{j=0}^{m-1} d(v, \partial K_j(v)) \quad \text{and} \quad 0 < \vartheta_m(v) = \frac{1}{m} \sum_{j=0}^{m-1} e^{-\beta d(v, \partial K_j(v))} \leq \vartheta$$

where $K_j(v) \in \mathscr{K}_j$ and we write $K(v)$ when $v \in K$ for some K in some $\mathscr{K}$. Here, θ is the smallest average distance between any site and the borders $\partial K_j(v) = \{v' \in K_j(v) : N(v') \not\subseteq K_j(v)\}$ of those blocks containing it, while ϑ captures a similar property in a more round about manner. We now define [14] the blocking operator for some partition $\mathscr{K}$

$$\mathsf{B}(\mathscr{K})\rho \triangleq \bigotimes_{K \in \mathscr{K}} \mathsf{B}^K \rho$$

where for any measure ρ on $\mathbb{X} = \bigotimes_{v \in V} \mathbb{X}^v$ and $J \subseteq V$ we denote by $\mathsf{B}^J \rho$ the marginal of ρ on $\bigotimes_{v \in J} \mathbb{X}^v$. The random field described by the measure $\mathsf{B}(\mathscr{K})\rho$ on $\mathbb{X}$ is independent across blocks defined by the partition $\mathscr{K}$.

The adaptively blocked particle filter adds a blocking operation into the bootstrap particle filter recursion

$$\hat{\pi}_0^\mu = \mu, \quad \hat{\pi}_n^\mu = \hat{\mathsf{F}}_n \hat{\pi}_{n-1}^\mu \quad (n \geq 1)$$

where $\hat{\mathsf{F}}_n \triangleq \mathsf{C}_n \mathsf{B}(\mathscr{K}_{\sigma(n)}) \mathsf{S}^N \mathsf{P}$ consists of four operations

$$\hat{\pi}_{n-1}^\mu \xrightarrow{\text{prediction/sampling}} \hat{\pi}_{n-}^\mu = \mathsf{S}^N \mathsf{P} \hat{\pi}_{n-1}^\mu \xrightarrow{\text{blocking/correction}} \hat{\pi}_n^\mu = \mathsf{C}_n \mathsf{B}(\mathscr{K}_{\sigma(n)}) \hat{\pi}_{n-}^\mu$$

where $\sigma : \mathbb{N} \to \{0, \ldots, m-1\}$ is a partition switching signal. If $m = 1$ then the adaptively blocked particle filter reduces to the blocked particle filter considered in [14]. If $\mathscr{K}_0 = \ldots = \mathscr{K}_{m-1} = \{V\}$ then the adaptively blocked particle filter reduces to the bootstrap particle filter. The resulting algorithm is given in Algorithm 1.

Algorithm 1 Adaptively Blocked Particle Filter

consider the partitions $\mathscr{K}_0, \ldots, \mathscr{K}_{m-1}$

let $\hat{\pi}_0^\mu = \mu$

for $k = 1, \ldots, n$ **do**

resample i.i.d. $\hat{x}_{k-1}(i) \sim \hat{\pi}_{k-1}^\mu$, $i = 1, \ldots, N$

sample $x_k^v(i) \sim p^v(\hat{x}_{k-1}(i), x^v)\psi(dx^v)$, $i = 1, \ldots, N$, $\forall v \in V$

compute $w_k^K(i) = \prod_{v \in K} g^v(x_k^v(i), Y_k^v) \Big/ \sum_{j=1}^N \prod_{v \in K} g^v(x_k^v(j), Y_k^v)$, $i = 1, \ldots, N$, $\forall K \in \mathscr{K}_{\sigma(k)}$

let $\hat{\pi}_k^\mu = \bigotimes_{K \in \mathscr{K}_{\sigma(k)}} \sum_{i=1}^N w_k^K(i)\, \delta_{x_k^K(i)}$

end for

The only difference between the adaptively blocked particle filter and the block particle filter of [14] is the adaptive consideration of multiple field partitions during the execution of the algorithm. Going forward, the notation $\hat{\pi}_n^\mu$ will refer to the adaptively blocked particle filter of Algorithm 1.

At any time n, only measurements in block $K \in \mathscr{K}_{\sigma(n)}$ are used to update the filter in block K. Each block $K \in \mathscr{K}_{\sigma(n)}$ can therefore be updated in parallel. The complexity of updating each block with a given error is thus dependent only on the cardinality of that block and not on the dimension of the entire random field. If the additional error for the entire filter caused by the blocking approximation can be sufficiently controlled, it seems the curse of dimensionality as it applies to the particle filter in general can be alleviated.

It is clear that at any given iteration the blocking operator decouples the distribution at the boundaries of the blocks. In the motivating paper [14], only a single partition is considered and it follows that the filtering approximation error will be larger at those vertices close to the boundary of each block than at those sites toward the centre of each block. The result is a spatially non-homogeneous filtering error and indeed the error bound derived in the motivating paper [14] is spatially dependent. A comprehensive and insightful discussion on this problem is provided in [14]. By adaptively applying different partitions one may ensure that, on average, each site of the random field is (at least approximately) the same distance from the borders of the blocks during some cycle. One would hope then that the error is, on average, spatially homogeneous and that one can achieve an error bound that is site independent. We note with this in mind the important special case of Algorithm 1 in which $\sigma(n) = n \,(\mathrm{mod}\, m)$ or where the partitions are applied in a cyclical order.

We stress that beyond the possibility of spatial smoothing, allowing for multiple partitions of the field may also provide algorithmic robustness in those cases in which the random field is time-varying since the partitions can be adapted online, or it may allow one to adaptively focus computation on certain locations of interest for periods of time etc. The design of σ offers seemingly much flexibility and other advantages of partition adaptation are envisioned but the discussion here will focus on spatial averaging of the blocking bias.

4 The Main Results

As in [14] we define the following norm

$$|\!|\!|\rho - \rho'|\!|\!|_J \triangleq \sup_{f\in\mathfrak{X}^J:|f|\le 1} \mathbf{E}\left[|\rho(f)-\rho'(f)|^2\right]^{1/2}$$

between two random measures ρ and ρ' on $\mathbb{X}$ where $\mathfrak{X}^J$ is the class of all measurable functions $f:\mathbb{X}\to\mathbb{R}$ with $f(x) = f(\tilde{x})$ whenever $x^J = (x^j)_{j\in J}$ for $J \subseteq V$. For example, one then finds $|\!|\!|\rho - \mathsf{S}^N\rho|\!|\!|_V \le 1/\sqrt{N}$ which captures the typical Monte Carlo approximation error. The goal is to bound the error between the nominal (ideal) Bayesian filter π_n^μ and the adaptively blocked particle filter $\hat{\pi}_n^\mu$. Recall that both the ideal filter and the adaptively blocked particle filter are defined recursively

$$\pi_n^\mu = \mathsf{F}_n\cdots\mathsf{F}_1\mu, \quad \hat{\pi}_n^\mu = \hat{\mathsf{F}}_n\cdots\hat{\mathsf{F}}_1\mu$$

where $\mathsf{F}_n \triangleq \mathsf{C}_n\mathsf{P}$ and $\hat{\mathsf{F}}_n \triangleq \mathsf{C}_n\mathsf{B}(\mathscr{K}_{\sigma(n)})\mathsf{S}^N\mathsf{P}$. From the triangle inequality we get

$$|\!|\!|\pi_n^\mu - \hat{\pi}_n^\mu|\!|\!|_J \le |\!|\!|\pi_n^\mu - \tilde{\pi}_n^\mu|\!|\!|_J + |\!|\!|\tilde{\pi}_n^\mu - \hat{\pi}_n^\mu|\!|\!|_J$$

for some $J\subseteq V$ where

$$\tilde{\pi}_n^\mu = \tilde{\mathsf{F}}_n\cdots\tilde{\mathsf{F}}_1\mu$$

with $\tilde{\mathsf{F}}_n \triangleq \mathsf{C}_n\mathsf{B}(\mathscr{K}_{\sigma(n)})\mathsf{P}$ is an ideal adaptively blocked filter; i.e. considering the adaptive blocking operation as it applies to the (non-sampled version of the) ideal Bayesian filter.

The expectation appearing in the definition of $|\!|\!|\cdot|\!|\!|_J$ is taken only with respect to the random sampling S^N; see [14]. Hence,

$$|\!|\!|\pi_n^\mu - \tilde{\pi}_n^\mu|\!|\!|_J = \sup_{f\in\mathfrak{X}^J:|f|\le 1} \mathbf{E}\left[|\pi_n^\mu(f)-\tilde{\pi}_n^\mu(f)|^2\right]^{1/2} = \sup_{f\in\mathfrak{X}^J:|f|\le 1} |\pi_n^\mu(f)-\tilde{\pi}_n^\mu(f)|$$

since no sampling occurs in π_n^μ or in $\tilde{\pi}_n^\mu$ and in this case $|\!|\!|\cdot|\!|\!|_J$ defines a local version of the total variation [10] which, as in [14], we sometimes denote by $\|\cdot\|_J$ for $J\subseteq V$.

The first term in this decomposition quantifies the bias introduced by the blocking operation alone. The second term quantifies the error due to the variance of the random sampling. Typical analysis on the convergence of the bootstrap particle filter deals only with a variance term (since $\tilde{\pi}_n^\mu = \pi_n^\mu$ in that case).

Let $\Delta \triangleq \max_{v\in V} \mathrm{card}\{v'\in V : d(v,v')\le r\}$ and $\Delta_{\mathscr{K}} \triangleq \max_s \max_{K\in\mathscr{K}_s} \mathrm{card}\{K'\in\mathscr{K}_s : d(K,K')\le r\}$. Also define $\Delta_d(v) \triangleq \max_s d(v,\partial K_s)$, $\Delta_d \triangleq \max_v \Delta_d(v)$ and $\nabla_d(v) \triangleq \min_s d(v,\partial K_s)$. Finally, let $|\mathscr{K}|_\infty \triangleq \max_s \max_{K\in\mathscr{K}_s}|K|$.

Theorem 1 (Bounding the Variance)**.** *There exists a constant $0<\varepsilon_0<1$, depending only on Δ and $\Delta_{\mathscr{K}}$ such that the following holds. Suppose there exist $\varepsilon_0<\varepsilon<1$ and $0<\kappa<1$ such that*

$$\varepsilon \le p^v(x,z^v) \le \varepsilon^{-1}, \quad \kappa \le g^v(x^v,y^v) \le \kappa^{-1} \quad \forall\, v\in V,\ x,z\in\mathbb{X},\ y\in\mathbb{Y}$$

Then for every $n\ge 0$, $x\in\mathbb{X}$, and $v\in V$ we have

$$|\!|\!|\tilde{\pi}_n^\mu - \hat{\pi}_n^\mu|\!|\!|_v \ \le\ \alpha\frac{|\mathscr{K}|_\infty e^{\beta|\mathscr{K}|_\infty}}{\sqrt{N}}$$

where $0<\alpha,\beta<\infty$ depend only on ε, κ, r, Δ and $\Delta_{\mathscr{K}}$.

The variance depends on the dimension of the sampling and so is necessarily dependent on the size of the blocks. This is exactly what is expected [8, 4] of the variance in the sense that it recovers the behaviour of the standard bootstrap particle filter (without blocking) which is dependent on the size of the entire field.

The blocking operation essentially reduces the large-scale filtering problem to one of multiple, smaller, independent filtering problems and on which each independent particle filter mostly exhibits an error that is well understood [8, 4]. Since the general nature (not the specific constants) of the variance bound is the best one might expect, we will not focus on this bound going forward.

The main error component of relevant interest here is that component introduced purely as a result of the blocking operation (and not the sampling). Therefore, going forward we are largely concerned with $\tilde{\pi}_n^\mu = \tilde{\mathsf{F}}_\mathsf{n} \cdots \tilde{\mathsf{F}}_1 \mu$ where $\tilde{\mathsf{F}}_n \triangleq \mathsf{C}_n \mathsf{B}(\mathscr{K}_{\sigma(n)})\mathsf{P}$ and its ability to approximate the ideal, full, Bayesian filter $\pi_n^\mu = \mathsf{F}_\mathsf{n} \cdots \mathsf{F}_1 \mu$. The particle filter $\hat{\pi}_n^\mu$ in this case can be thought of as an approximation of the ideal blocked filter $\tilde{\pi}_n^\mu$ and it is worth noting that other approximations to $\tilde{\pi}_n^\mu$ separate to particle-based approximations could be substituted.

In summary, the main contribution is reduced to a study on adaptively blocked filtering and the error introduced through adaptive blocking when compared to the ideal Bayesian filter. The particle filtering step is given to show how one may approximate the adaptively blocked filter in practice and the error given on this particle representation is noted for completeness (this error is as expected even if it is non-trivial to derive).

The bound on the bias introduced due to blocking is now stated.

Theorem 2 (Bounding the Bias)**.** *Suppose that $\varepsilon \leq p^\nu(x, z^\nu) \leq \varepsilon^{-1}$ for all $\nu \in V$ and $x, z \in \mathbb{X}$ with $\varepsilon > \varepsilon_0 = \left(1 - 1/(18\Delta^2)\right)^{1/2\Delta}$. Let $\beta = -(2r)^{-1} \log 18\Delta^2(1-\varepsilon^{2\Delta}) > 0$. If $\sigma(s) = s \pmod m$, $s \in \mathbb{N}$ then for every $\nu \in V$ we have*

$$\frac{1}{m}\sum_{k=0}^{m-1} \|\pi_{n-k}^x - \tilde{\pi}_{n-k}^x\|_\nu \;\leq\; \frac{8e^{-\beta}}{(1-e^{-\beta})}(1-\varepsilon^{2\Delta})\vartheta_m(\nu) \;\leq\; \frac{8e^{-\beta}}{(1-e^{-\beta})}(1-\varepsilon^{2\Delta}) \exp\left[-\beta e^{-\beta(\Delta_d(\nu)-\nabla_d(\nu))}\frac{1}{m}\theta_m(\nu)\right]$$

for every $n \geq 0$ and $x \in \mathbb{X}$ where $\theta_m(\nu) = \frac{1}{m}\sum_{j=0}^{m-1} d(\nu, \partial K_j(\nu))$ and $\vartheta_m(\nu) = \frac{1}{m}\sum_{j=0}^{m-1} e^{-\beta d(\nu, \partial K_j(\nu))}$ and $K_j(\nu) \in \mathscr{K}_j$.

This is a time-uniform bound on the average bias over a time length of m. Both inequalities in Theorem 2 imply that the bias introduced due to blocking can be spatially averaged (smoothed) across a cyclical application of a sequence of partitions. Both inequalities collapse to the result of [14] in the case $m = 1$. The second inequality, in general, over bounds the first inequality but may be more convenient for discussion as $\theta_m(\nu)$ may be easier than $\vartheta_m(\nu)$ to conceptualise. Following the analysis in [14], the goal was to derive a similarly natured bound here, but which captured honestly the spatial smoothing effect.

The spatial invariance of the error bound (or more specifically the bias bound) will be discussed in more detail in the next section. We simply note here that if $\theta = \theta_m(\nu) = \frac{1}{m}\sum_{j=0}^{m-1} d(\nu, \partial K_j(\nu))$ where $K_j(\nu) \in \mathscr{K}_j$ for all $\nu \in \mathscr{V}$, then the bound really is spatially invariant. Such a situation occurs for a particular class of graphs (i.e. random fields) discussed later. The more general case in which $\theta \leq \frac{1}{m}\sum_{j=0}^{m-1} d(\nu, \partial K_j(\nu))$ is also discussed.

A simple corollary follows in which there exists an ordering of $\mathscr{K}_0, \ldots, \mathscr{K}_{m-1}$ such that for a cyclical sequence of partitions $\sigma(s) = s \pmod m$, $s \in \mathbb{N}$ we have

$$\|\pi_{n-k}^x - \tilde{\pi}_{n-k}^x\|_\nu \;\leq\; \tfrac{8e^{-\beta}}{(1-e^{-\beta})}(1-\varepsilon^{2\Delta})\vartheta_m(\nu) \;\leq\; \tfrac{8e^{-\beta}}{(1-e^{-\beta})}(1-\varepsilon^{2\Delta}) \exp\left[-\beta e^{-\beta(\Delta_d(\nu)-\nabla_d(\nu))}\frac{1}{m}\theta_m(\nu)\right]$$

for some (at least one) $\nu \in V$.

The blocking operation contributes the bias term to the total error while the random sampling contributes the variance term. As previously noted, the nature of the variance bound is as expected and we do not focus on that going forward. The bias is determined by the blocking operator which conceptually, at any time, has little effect on those sites far removed from the block borders due to the local dynamical dependencies assumed. Consequently, the bias is controlled at a sub-block level in that the bias at each site is controlled, on average, by that site's distance to the border of the blocks which contain it. If we can average this distance across the field through adaptive partitioning then we should be averaging the bias across the field.

From the computational view point, the variance bound implies that one need only consider the size of the blocks (not the entire random field) when picking a value for N to control the error. This (ideally) leads to a reduction in the computational requirements of the filtering problem and is the underlying motivation for blocking. This gain comes at a price, in the form of a bias introduced due to blocking. Here we will consider adaptive blocking as a way of smoothing the bias error over the random field or controlling the bias in a more precise way.

Finally, we refer to [14] for a discussion on the mixing assumption $\varepsilon \leq p^\nu(x, z^\nu) \leq \varepsilon^{-1}$ with the non-standard requirement $\varepsilon_0 < \varepsilon < 1$ with $\varepsilon_0 > 0$. The restrictions on $\varepsilon_0 < \varepsilon < 1$ are relaxed partially in [15]. This term does not alter the significance of the result at the proof-of-concept stage. Moreover, non-optimal restrictions on the system model in this form appear frequently in similar studies; e.g. see [8]. Typically, the empirical evidence suggests a far more relaxed application of the algorithms in question is permissible.

4.1 Proof Strategy

The bias and variance bounds are treated separately but lead to a total error bound. The proof strategy is adopted from Rebeschini and van Handel [14] and much of the analysis required is identical and not repeated. The strategy in [14] is inspired in part by the time-uniform particle filter convergence results [7].

In the case of the bias $\|\pi_n^\mu - \tilde{\pi}_n^\mu\|_J$, one first derives a local stability property for the filter π_n^μ which implies that the marginal over a local set $J \subseteq V$ of the initial state μ is forgotten exponentially fast. Such a property also implies that any approximation errors in, say, the initial state are also forgotten. It then follows that if one can bound the one-step approximation error $\|\mathsf{F}_n\tilde{\pi}_{n-1}^\mu - \tilde{\mathsf{F}}_n\tilde{\pi}_{n-1}^\mu\|_J$ at any time, then in conjunction with the local stability property one will obtain a time-uniform bound on the bias over a local region of the field.

In the case of the variance $\|\tilde{\pi}_n^\mu - \hat{\pi}_n^\mu\|_J$, a similar idea is used except one first establishes stability for the ideal adaptively blocked filter $\tilde{\pi}_n^\mu$. Then, one must bound the one-step approximation error $\|\tilde{\mathsf{F}}_n\hat{\pi}_{n-1}^\mu - \hat{\mathsf{F}}_n\hat{\pi}_{n-1}^\mu\|_J$ at any time. Putting the stability property and the bound on the one-step approximation together, one achieves the desired time-uniform bound on the variance of a block in the adaptively blocked filter.

We have obviously glossed over much of the intricacies involved in the proof in this summary. For example, in the case of the bias, the property introduced in [14] and referred to as the decay of correlations must be established to hold uniformly in time for the ideal block filter $\tilde{\pi}_n^\mu$. This property captures a notion of spatial stability where the state at some site in the random field is forgotten as one moves away from that site. Rebeschini et al. provide a novel measure of this decay that allows them to establish local stability of the filter π_n^μ and to establish a bound on the one-step approximation error $\|\mathsf{F}_n\tilde{\pi}_{n-1}^\mu - \tilde{\mathsf{F}}_n\tilde{\pi}_{n-1}^\mu\|_J$. Conceptually, a property like the decay of correlations is necessary to establish such results.

We refer the reader to [14] for a broader, more insightful, discussion on the strategy. Now, we note specifically what ideas must be altered to account for a change in partition from one time to the next.

4.1.1 Steps to Prove the Bias Bound

Very roughly speaking the steps needed to prove the bound on the bias include: 1). Establishing the local stability property for the nonlinear filter π_n^μ; and, 2). establishing the decay of correlations property holds uniformly in time for the ideal adaptively blocked filter $\tilde{\pi}_n^\mu$; and then, 3). establishing a bound on the one-step approximation error $\|\mathsf{F}_n\tilde{\pi}_{n-1}^\mu - \tilde{\mathsf{F}}_n\tilde{\pi}_{n-1}^\mu\|_J$ that holds at any time; and finally, 4) putting it all together.

The local stability of π_n^μ depends on the decay of correlations property assumed on μ and is otherwise independent of the blocking procedure. Hence, we can take this result as a given [14]. The one-step approximation error $\|\mathsf{F}_n\tilde{\pi}_{n-1}^\mu - \tilde{\mathsf{F}}_n\tilde{\pi}_{n-1}^\mu\|_J$ is dependent on the blocking procedure but only on the partition in effect during a single time step, and this partition is otherwise arbitrary, so we can take this result as given [14].

To prove our case, we only need to establish that the decay of correlations property holds uniformly in time for the filter distribution $\tilde{\pi}_n^\mu$ when given the changing partitions. Once this is established, it is just a matter of collecting the relevant results and finalising the bound on $\|\pi_n^\mu - \tilde{\pi}_n^\mu\|_J$. We follow through with this last step and show how the spatial averaging effect of $\theta_m(v) = \frac{1}{m}\sum_{j=0}^m d(v, \partial K_j(v))$ comes out during this procedure.

The detailed proof is given in a subsequent section drawing from [14] as often as possible.

4.1.2 Steps to Prove the Variance Bound

Roughly again, the steps needed to prove the bound on the variance include: 1). Establishing the stability of the ideal adaptively blocked filter $\tilde{\pi}_n^\mu$; and, 2). establishing a bound on the one-step approximation error $\|\tilde{\mathsf{F}}_n\hat{\pi}_{n-1}^\mu - \hat{\mathsf{F}}_n\hat{\pi}_{n-1}^\mu\|_J$ that holds at any time; and finally, 3). putting it all together.

Firstly, we do not have to deal with any correlation-like properties in the case of the variance bound and as noted the final result is as expected. So things may appear simpler initially. Unfortunately, proving stability for the ideal adaptively blocked filter $\tilde{\pi}_n^\mu$ is not trivial [14]. Because the stability of $\tilde{\pi}_n^\mu$ is dependent on the change of partition we must re-establish that this stability result holds in the case of adaptively changing partitions for completeness. Moreover, for technical reasons related to the use of the norm $\|\cdot\|_J$, the authors in [14] consider instead a two-step approximation error $\|\tilde{\mathsf{F}}_{n+1}\tilde{\mathsf{F}}_n\hat{\pi}_{n-1}^\mu - \tilde{\mathsf{F}}_{n+1}\hat{\mathsf{F}}_n\hat{\pi}_{n-1}^\mu\|_J$ and bound this term at any time. Because a change in partition comes into play over two steps, we must re-establish that this two-step approximation error is bounded under a change of partition at any time. We then bring the relevant results together and finalise the bound on $\|\tilde{\pi}_n^\mu - \hat{\pi}_n^\mu\|_J$.

As noted, this is the strategy taken to derive the variance bound but, since our main concern here is the spatial aspects of the filtering problem and the related (adaptive) blocking operation, we do not give the details. Many of the technical lemmas involved in this analysis also follow directly from [14] and those results requiring a modification to their proofs need only an arguably minor re-analysis and modification. The final result is as expected and the authors are available to provide the variance bound proof details on request.

5 Discussion on the Adaptively Blocked Filter and Spatial Smoothing

The main result in the previous section is a total error bound on $\|\!|\pi_n^\mu - \hat{\pi}_n^\mu|\!\|_v$, for all $v \in V$ and can be decomposed (and is actually derived) in terms of a bound on the variance (induced by the random Monte Carlo procedure) and a bound on the bias (induced by the blocking operator). We focus on the bias bound in this section and its relevance as it pertains to the dependence of the total error on the particular spatial site v.

The main point of interest in this work is the effect of the (adaptive) blocking operation on the total error bound which shows up purely via the systematic bias. To this end, we compare the bound on the bias proposed here with the bias bound proposed in the motivating paper [14] by Rebeschini et al.,

$$\text{bias}_{\text{B}}(v) \leq \mathcal{O}(\tfrac{1}{m}\textstyle\sum_{j=0}^{m-1} e^{-d(v,\partial K_j(v))}) = \mathcal{O}(\vartheta_m(v)) \leq \mathcal{O}(e^{\frac{-1}{m}\sum_{j=0}^{m-1} d(v,\partial K_j(v))}) = \mathcal{O}(e^{-\theta_m(v)}) \qquad \text{Bertoli et al.}$$

$$\text{bias}_{\text{R}}(v) \leq \mathcal{O}(e^{-d(v,\partial K(v))}) \qquad \text{Rebeschini et al.}$$

where Rebeschini et al. only ever consider a single partition. Here, $\text{bias}_{\text{B}}(v)$ and $\text{bias}_{\text{R}}(v)$ can be taken as the average of the bias over a time period of length m. We remove any unnecessary constants from the expressions that cloud the conceptual discussion. Here, we use $\mathcal{O}(\cdot)$ to capture only that spatially dependent component of the bias bound noting that the constant 'out-the-front' is equivalent in both cases [14]. For conceptual, rather than technical, reasons we consider the slightly looser bound $\text{bias}_{\text{B}}(v) \leq \mathcal{O}(e^{-\theta_m(v)})$ during discussion.

We highlight that if $\theta_m(v) = \frac{1}{m}\sum_{j=0}^{m-1} d(v,\partial K_j(v))$ for all $v \in \mathcal{V}$ then

$$\mathcal{O}(e^{\frac{-1}{m}\sum_{j=0}^{m-1} d(v,\partial K_j(v))}) = \mathcal{O}(e^{-\theta_m(v)}) = \mathcal{O}(e^{-\theta})$$

and the bound is truly spatially invariant. This is part of the motivation for this work and is explored in more detail now. Consider Figure 1.

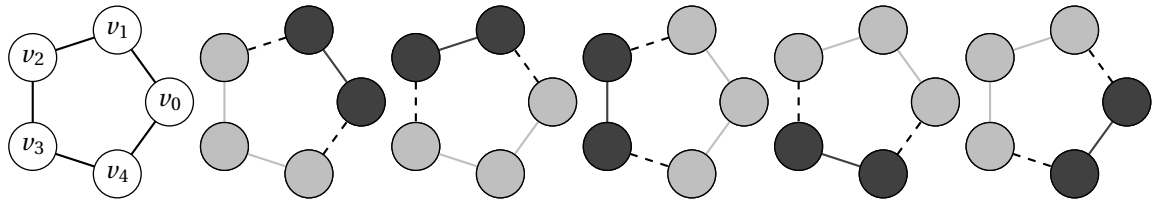

Figure 1: The left most circulant graph depicts the underlying graphical structure of the random field. The remaining five graphs from second left to far right highlight a sequence of independent partitions of the original graph in which each partition consists of a set of two (dark grey) and three (light grey) vertices.

Pick any site $v \in V$ in the graph depicted in Figure 1 and note that $\theta = \theta_m(v) = 1/5$. One then clearly has a spatially uniform bound on $\text{bias}_{\text{B}}(v)$. Consider now any single partition alone and note that for four out of the five sites we have $d(v,\partial K(v)) = 0$ and at one site we have $d(v,\partial K(v)) = 1$ which implies, as noted by Rebeschini et al., that the bound on $\text{bias}_{\text{R}}(v)$ is not spatially uniform. The adaptive blocking procedure is averaging the distance $d(v,\partial K_j(v))$ through the use of multiple partitions which results in a kind of spatial error smoothing.

The bound on the bias of the cyclically blocked filter at every site $v \in V$ is completely independent of the site v in every case in which $\theta = \theta_m(v)$, $\forall v \in V$. Such cases may occur in practice; e.g. a sequence of partitions on any regular lattice wrapped on a torus can be derived that obeys this property, see Figure 2.

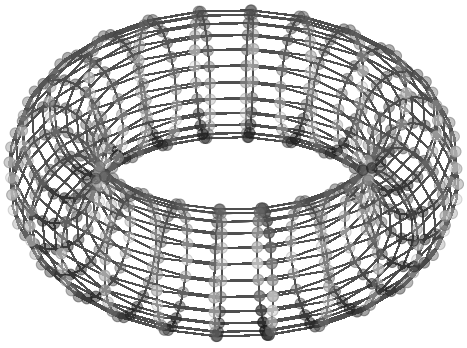

Figure 2: A lattice wrapped on a torus with no physical boundaries.

In the general case, in which $\theta \leq \theta_m(v)$ for some $v \in V$, the spatial smoothing property of the cyclical blocking filter is still in effect and reduces, as compared to [14], the degree of spatial inhomogeneity (on average) as it applies to the bias bound. Essentially, the sites on the borders of a block in $\mathscr{K}$ are typically not on the borders of a block in $\mathscr{K}'$, while the sites at the centre of a block in $\mathscr{K}$ are typically not at the centre of a block in $\mathscr{K}'$. Given a sufficient number of well-chosen partitions of this type, then one can ensure the average distance of a site to a border is smoothed (or spatially averaged) across all sites. Of course, this means that some particular sites may be worse off than they were under a single partition (this is an obvious consequence averaging).

For example, consider again the case in Figure 1 but suppose only the four left most partitions are employed by the cyclically blocked particle filter. Then $\theta_m(v_0) = \theta_m(v_1) = \theta_m(v_3) = \theta_m(v_4) = 1/4$ while $\theta_m(v_2) = 0$. Clearly, one has a more desirable bound on the bias on average in this case than in the case in which only a single partition is considered, albeit complete spatial homogeneity is not achieved. Considering additional partitions in a large-scale random field will be of even further benefit than that exposited in this toy example.

Note finally that in both Rebeschini et al. and here the spatial uniformity of the error (or the bias more specifically) is often referred to via the bound on the bias and not the bias or error itself. This is a consequence of the technical analysis, but for all practical purposes it would appear obvious that the spatial homogeneity of the error itself is of the same nature as that noted by the bound applicable to that error (even if such bounds are otherwise quite loose). That is, the blocked filter of Rebeschini et al. [14] would clearly seem to favour those sites far from the border of the individual blocks in terms of the actual performance of the filter, while the cyclically blocked filter proposed in this work is clearly, in some sense, averaging out this favouritism and its effect on the actual filter performance at any site. The point is that the spatial relationship of the error is typically noted in terms of the bias bounds but intuitively/conceptually the discussions on homogeneity (or inhomogeneity) of a particular filter apply (seemingly) also to the error/bias itself.

6 Proof of the Bias Bound

Rebeschini et al. [14] introduced an important concept referred to as the decay of correlation which captures, in a very technical manner, the intuitive notion of spatial stability where the state at some site in the random field should be forgotten as one moves away from that site. This notion plays a crucial role in the convergence of the bias due to the blocking operation.

It is important to note that the analysis and the spatial stability property put forth in [14] is based in part on those ideas of temporal stability introduced in [8] and used to establish time-uniform convergence results for the standard bootstrap particle filter.

Recall that the dynamics of the underlying process $(X_n)_{n\geq 0}$ are local in the sense that $p^v(x, z^v)$ depends only on $x^{N(v)}$ where $x^J = (x^j)_{j\in J}$ for $J \subseteq V$. Here, $N(v) = \{v' \in V : d(v, v') \leq r\}$ for some $r \in \mathbb{N}$ captures the local neighbourhood of sites on which site v explicitly depends.

We now briefly review the measure introduced in [14] on the decay of correlations property. For any probability measure μ on $\mathbb{X}$ and for $x, z \in \mathbb{X}$ with $v \in V$ define

$$\mu^v_{x,z}(A) \triangleq \mathbf{P}^\mu[X^v_{n-1} \in A | X^{V\setminus\{v\}}_{n-1} = x^{V\setminus\{v\}},\ X_n = z] = \frac{\int \mathbf{1}_A(x^v) \prod_{u\in N(v)} p^u(x, z^u) \mu^v_x(dx^v)}{\int \prod_{u\in N(v)} p^u(x, z^u) \mu^v_x(dx^v)}$$

for any $A \in \mathscr{X}$ and where

$$\mu^v_x(dx^v) \triangleq \mu(X^v_{n-1} = dx^v | X^{V\setminus\{v\}}_{n-1} = x^{V\setminus\{v\}})$$

Then

$$C^\mu_{vv'} \triangleq \frac{1}{2} \sup_{z\in\mathbb{X}} \sup_{x,\tilde{x}\in\mathbb{X}\,:\, x^{V\setminus\{v'\}} = \tilde{x}^{V\setminus\{v'\}}} \|\mu^v_{x,z} - \mu^v_{\tilde{x},z}\|$$

for $v, v' \in V$. This quantity $C^\mu_{vv'}$ somehow captures the correlation between two sites $v, v' \in V$ in the random field under the assumed field model. A little more precisely, this term is measuring the maximal total variation at a site v that may arise due to a perturbation at site v'. Now define

$$\mathrm{Corr}(\mu, \beta) \triangleq \max_{v\in V} \sum_{v'\in V} e^{\beta d(v,v')} C^\mu_{vv'}$$

with $\beta > 0$. This quantity $\mathrm{Corr}(\mu, \beta)$ is a measure on the total degree of correlation decay for the measure μ given a rate parameter β. The site v can be interpreted as the most sensitive site in the field. To understand this quantity $\mathrm{Corr}(\mu, \beta)$ a little more conceptually, suppose the most sensitive site v is known a priori. Then suppose that $\sum_{v'\in V} e^{\beta d(v,v')} C^\mu_{vv'} = 1$. It follows that the correlation between any two sites in the random field decays as a function of the distance between those two sites at an exponential rate defined by at least β.

With $\mathrm{Corr}(\mu,\beta)$ defined as such we borrow directly from [14] the local stability result on π_n^μ which requires only that the initial condition μ satisfy a decay of correlations property. Recall that $\Delta \triangleq \max_{v\in V} \mathrm{card}\{v' \in V : d(v,v') \le r\}$ defines the size of the largest neighbourhood in V.

Lemma 1 (Local Filter Stability [14])**.** *Suppose there exists $\varepsilon > 0$ such that $\varepsilon \le p^v(x,z^v) \le \varepsilon^{-1}$ for all $v \in V$ and $x,z \in \mathbb{X}$. Let μ,ν be probability measures on $\mathbb{X}$, and suppose that*

$$\mathrm{Corr}(\mu,\beta) \le \frac{1}{2} - 3(1-\varepsilon^{2\Delta})e^{2\beta r}\Delta^2$$

for a sufficiently small constant $\beta > 0$. Then

$$|||\mathsf{F}_n \cdots \mathsf{F}_{s+1}\mu - \mathsf{F}_n \cdots \mathsf{F}_{s+1}\nu|||_J \le 2e^{-\beta(n-s)} \sum_{v\in J} \max_{v'\in V} e^{-\beta d(v,v')} \sup_{x,z\in \mathrm{X}} |||\mu_{x,z}^{v'} - \nu_{x,z}^{v'}|||$$

for every $J \subseteq V$ and $s < n$.

Now consider any partition $\mathscr{K}_i$. We define a correlation depending on the given partition. Fix a probability measure μ on $\mathbb{X}$ and $x,z \in \mathbb{X}$, $v \in V$, $K \in \mathscr{K}_i$ and then let

$$\mu_{x,z}^{v,K}(A) \triangleq \mathbf{P}^\mu[X_{n-1}^v \in A | X_{n-1}^{V\setminus\{v\}} = x^{V\setminus\{v\}}, X_n^K = z^K] = \frac{\int \mathbf{1}_A(x^v) \prod_{u\in N(v)\cap K} p^u(x,z^u)\mu_x^v(dx^v)}{\int \prod_{u\in N(v)\cap K} p^u(x,z^u)\mu_x^v(dx^v)}$$

for any $A \in \mathscr{X}$. Now define

$$\widetilde{C}_{vv'}^{\mathscr{K}_i,\mu} \triangleq \frac{1}{2} \max_{K\in\mathscr{K}_i} \sup_{z\in\mathbb{X}} \sup_{x,\widetilde{x}\in\mathbb{X} \,:\, x^{V\setminus\{v'\}} = \widetilde{x}^{V\setminus\{v'\}}} \|\mu_{x,z}^{v,K} - \mu_{\widetilde{x},z}^{v,K}\|$$

and

$$\widetilde{\mathrm{Corr}}_{\mathscr{K}_i}(\mu,\beta) \triangleq \max_{v\in V} \sum_{v'\in V} e^{\beta d(v,v')} \widetilde{C}_{vv'}^{\mathscr{K}_i,\mu}$$

for any partition $\mathscr{K}_i$ in the cyclically blocked particle filter sequence. One can interpret the block adapted measure $\widetilde{\mathrm{Corr}}_{\mathscr{K}_i}(\mu,\beta)$ in much the same way as $\mathrm{Corr}(\mu,\beta)$.

The reason behind formulating the measure $\mathrm{Corr}(\mu,\beta)$ in such a way is technical and follows from the program put forth in [14]. Here we are only modifying this program to account for changing partitions and in doing so we seek to draw on the detailed analysis put forth in [14] as much as possible. The reason for introducing $\widetilde{\mathrm{Corr}}_{\mathscr{K}_i}(\mu,\beta)$ is now explained. In order to establish the one-step approximation error we must bound $\mathrm{Corr}(\tilde{\pi}_n^\mu,\beta)$ uniformly in time. The authors of [14] note the difficulty in working directly with $\mathrm{Corr}(\tilde{\pi}_n^\mu,\beta)$ and instead propose to bound $\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(n)}}(\tilde{\pi}_n^\mu,\beta)$ and then use the following result to indirectly control $\mathrm{Corr}(\tilde{\pi}_n^\mu,\beta)$.

Lemma 2 ([14])**.** *For any probability measure ν, partition $\mathscr{K}_i$ and $\beta > 0$, we have*

$$\mathrm{Corr}(\nu,\beta) \le (1-\varepsilon^{2\Delta})e^{2\beta r}\Delta^2 + 2\varepsilon^{-2\Delta}\widetilde{\mathrm{Corr}}_{\mathscr{K}_i}(\nu,\beta)$$

Given this result it follows that a time-uniform bound on $\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(n)}}(\tilde{\pi}_n^\mu,\beta)$ leads easily to a time-uniform bound on $\mathrm{Corr}(\tilde{\pi}_n^\mu,\beta)$.

Because $\tilde{\pi}_n^\mu$ is dependent on the changing partitions, which is central to the adaptively blocked filter, it follows that we must re-establish that a time-uniform bound on $\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(n)}}(\tilde{\pi}_n^\mu,\beta)$ exists in the case of interest here.

We prove a technical lemma that will be used subsequently.

Lemma 3. *For any strictly positive $a,b \in \mathbb{R}_+$ and $x \in \mathbb{R}_+$ with $x \ge 1$ the following holds: $|ax - b/x| \le |a-b|x^2$.*

Proof. If $a \ge b$ or even $ax \ge b/x$ then

$$|ax - \frac{b}{x}| \;=\; ax - bx + bx - \frac{b}{x} \;=\; |a-b|x + b\frac{x^2-1}{x} \;\le\; |a-b|x \;\le\; |a-b|x^2$$

Note $ax \ge b/x$ implies $x \ge \sqrt{b/a}$ for any $a,b,x > 0$. This leaves the case $b > a$ and $1 \le x < \sqrt{b/a}$. In this case

$$|ax - \frac{b}{x}| \;=\; |ax^2 - b|\frac{1}{x} \;\le\; |ax^2 - b| \;\le |a-b| \le |a-b|x^2$$

by continuity in $1 \le x < \sqrt{b/a}$ for any fixed $b > a > 0$. □

The following proposition relates $\widetilde{\mathrm{Corr}}_{\mathscr{K}_i}(\mu,\beta)$ under one partition to the same quantity considered under a different partition and is needed to obtain the time-uniform bound on the $\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(n)}}(\tilde{\pi}_n^\mu,\beta)$.

Proposition 1. *Suppose there exists $\varepsilon>0$ such that $\varepsilon \le p^v(x,z^v) \le \varepsilon^{-1}$. For any probability measure μ, partitions $\mathscr{K}_i,\mathscr{K}_j$ and $\beta>0$, we have*

$$\widetilde{\mathrm{Corr}}_{\mathscr{K}_i}(\mu,\beta) \le \varepsilon^{-8\Delta}\widetilde{\mathrm{Corr}}_{\mathscr{K}_j}(\mu,\beta)$$

Proof. Pick $K_i \in \mathscr{K}_i$ and $K_j \in \mathscr{K}_j$ with $K \triangleq K_i \cap K_j \neq \emptyset$. Then by definition

$$\begin{aligned}
\mu_{x,z}^{v,K_i}(A) &= \frac{\int \mathbf{1}_A(x^v)\prod_{u\in N(v)\cap K_i\setminus K} p^u(x,z^u)\mu_{x,z}^{v,K}(dx^v)}{\int \prod_{u\in N(v)\cap K_i\setminus K} p^u(x,z^u)\mu_{x,z}^{v,K}(dx^v)} \le \frac{\int \mathbf{1}_A(x^v)\varepsilon^{-|N(V)\cap K_i\setminus K|}\prod_{u\in N(v)\cap K_j\setminus K}\frac{p^u(x,z^u)}{\varepsilon}\mu_{x,z}^{v,K}(dx^v)}{\int \varepsilon^{|N(V)\cap K_i\setminus K|}\prod_{u\in N(v)\cap K_j\setminus K}\frac{p^u(x,z^u)}{\varepsilon^{-1}}\mu_{x,z}^{v,K}(dx^v)} \\
&\le \varepsilon^{-2[|N(V)\cap K_i\setminus K|+|N(V)\cap K_j\setminus K|]}\mu_{x,z}^{v,K_j}(A) \\
&\le \varepsilon^{-4\Delta}\mu_{x,z}^{v,K_j}(A)
\end{aligned}$$

for $A\in\mathscr{X}$. Alternatively,

$$\mu_{x,z}^{v,K_i}(A) \ge \varepsilon^{4\Delta}\mu_{x,z}^{v,K_j}(A)$$

noting that $\mathscr{K}_i$ and $\mathscr{K}_j$ are anyway arbitrary. Fix $A\in\mathscr{X}$ and $x,\tilde{x}\in\mathbb{X}$ and suppose, without loss of generality, $\mu_{x,z}^{v,K_i}(A) \ge \mu_{\tilde{x},z}^{v,K_i}(A)$. It follows

$$|\mu_{x,z}^{v,K_1}(A)-\mu_{\tilde{x},z}^{v,K_1}(A)| = \mu_{x,z}^{v,K_1}(A)-\mu_{\tilde{x},z}^{v,K_1}(A) \le \varepsilon^{-4\Delta}\mu_{x,z}^{v,K_2}(A)-\varepsilon^{4\Delta}\mu_{\tilde{x},z}^{v,K_2}(A) \le \varepsilon^{-8\Delta}|\mu_{x,z}^{v,K_2}(A)-\mu_{\tilde{x},z}^{v,K_2}(A)|$$

using Lemma 3. Taking the supremum over $A\in\mathscr{X}$ gives

$$\|\mu_{x,z}^{v,K_1}-\mu_{\tilde{x},z}^{v,K_1}\| \le \varepsilon^{-8\Delta}\|\mu_{x,z}^{v,K_2}-\mu_{\tilde{x},z}^{v,K_2}\|$$

and from this the result of the proposition follow easily. □

We state the following lemma which bounds the change in the correlation decay over a one-step application of the ideal cyclically blocked filter for a fixed partition.

Lemma 4 ([14])**.** *Suppose there exists $\varepsilon>0$ such that $\varepsilon \le p^v(x,z^v) \le \varepsilon^{-1}$ for all $v\in V$ and $x,z\in\mathbb{X}$. For any probability measure ν, partition $\mathscr{K}_{\sigma(s)}$ and sufficiently small $\beta>0$ such that*

$$\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(s)}}(\nu,\beta) \;\le\; 1/2-(1-\varepsilon^2)e^{\beta(r+1)}\Delta$$

then

$$\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(s)}}(\tilde{\mathsf{F}}_s\nu,\beta) \;\le\; 2(1-\varepsilon^{2\Delta})e^{2\beta r}\Delta^2$$

for any $s\in\mathbb{N}$.

Now we are in a position to prove the time-uniform bound on $\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(n)}}(\tilde{\pi}_n^\mu,\beta)$. We have via Lemma 4 a bound on the change in the correlation decay over any single time step. We have via Proposition 1 a relationship between the decay of correlations for a measure under two different partitions. We combine these results and iterate to get the desired time-uniform bound as now shown.

Lemma 5. *Assume $\varepsilon \le p^v(x,z^v) \le \varepsilon^{-1}$ for all $v\in V$ and $x,z\in\mathbb{X}$ with*

$$\varepsilon > \varepsilon_0 = \left(1-1/(16\Delta^2)\right)^{\frac{1}{2\Delta}}$$

Let μ be a probability measure on $\mathbb{X}$ and $\mathscr{K}_0$ a partition of V such that

$$\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(0)}}(\mu,\beta) \;\le\; 1/8$$

where $\beta = -\frac{1}{2r}\log[16\Delta^2(1-\varepsilon^{2\Delta})]>0$. Then

$$\widetilde{\mathrm{Corr}}_{\mathscr{K}_{\sigma(n)}}(\tilde{\pi}_n^\mu,\beta) \;\le\; 1/8$$

for all $n\ge 0$.

Proof. First, $\varepsilon > \varepsilon_0 = \left(1 - 1/(16\Delta^2)\right)^{\frac{1}{2\Delta}}$ implies $\varepsilon^{-8\Delta} \leq 2$ and $(1-\varepsilon^2)e^{\beta(r+1)}\Delta \leq 1/16$. Thus, given $\widetilde{\mathrm{Corr}}_{\mathcal{K}_{\sigma(0)}}(\mu,\beta) \leq 1/8$ we have

$$\widetilde{\mathrm{Corr}}_{\mathcal{K}_{\sigma(1)}}(\mu,\beta) \;\leq\; 2\widetilde{\mathrm{Corr}}_{\mathcal{K}_{\sigma(0)}}(\mu,\beta) \leq \frac{1}{4}$$

via Proposition 1. Now, given $\widetilde{\mathrm{Corr}}_{\mathcal{K}_{\sigma(1)}}(\mu,\beta) \leq 1/4 < 7/16 \leq 1/2 - (1-\varepsilon^2)e^{\beta(r+1)}$ we have

$$\widetilde{\mathrm{Corr}}_{\mathcal{K}_{\sigma(1)}}(\tilde{\pi}_1^{\mu},\beta) \;=\; \widetilde{\mathrm{Corr}}_{\mathcal{K}_{\sigma(1)}}(\tilde{\mathsf{F}}_1\mu,\beta) \;\leq\; 2(1-\varepsilon^{2\Delta})e^{2\beta r}\Delta^2 \;\leq\; \frac{1}{8}$$

via Lemma 4. Just restart the argument and iterate to get $\widetilde{\mathrm{Corr}}_{\mathcal{K}_{\sigma(n)}}(\tilde{\pi}_n^{\mu},\beta) \leq 1/8$ for any $n \geq 0$. □

Now we have a time-uniform bound on $\widetilde{\mathrm{Corr}}_{\mathcal{K}_{\sigma(n)}}(\tilde{\pi}_n^{\mu},\beta)$, and we can use Lemma 2 to arrive at the result we want which is a time-uniform bound on the decay of correlation measure of interest.

Lemma 6 (Bounding the Decay of Correlations)**.** *Assume $\varepsilon \leq p^{v}(x,z^{v}) \leq \varepsilon^{-1}$ for all $v \in V$ and $x,z \in \mathbb{X}$ with*

$$\varepsilon > \varepsilon_0 = \left(1 - 1/(16\Delta^2)\right)^{\frac{1}{2\Delta}}$$

Let $\beta = -(2r)^{-1}\log 16\Delta^2(1-\varepsilon^{2\Delta}) > 0$. Then

$$\mathrm{Corr}(\tilde{\pi}_n^{\mu},\beta) \leq \frac{1}{3}$$

for every $n \geq 0$ and any probability measure μ on $\mathbb{X}$.

Recall the program required to prove the bias bound: 1). Establish the local stability property for the non-linear filter π_n^{μ}; and, 2). establish the decay of correlations property holds uniformly in time for the ideal cyclical blocked filter $\tilde{\pi}_n^{\mu}$; and then, 3). establish a bound on the one-step approximation error $|||\mathsf{F}_n\tilde{\pi}_{n-1}^{\mu} - \tilde{\mathsf{F}}_n\tilde{\pi}_{n-1}^{\mu}|||_J$ that holds at any time; and finally, 4). put it all together.

All that remains is to show that $|||\mathsf{F}_n\tilde{\pi}_{n-1}^{\mu} - \tilde{\mathsf{F}}_n\tilde{\pi}_{n-1}^{\mu}|||_J$ is bounded at any time and then to put it all together. Since this one-step approximation error is independent of any change in partition we refer back to [14].

Lemma 7 (Bounding the One-Step Blocking Approximation Error [14])**.** *Suppose there exists $\varepsilon > 0$ such that $\varepsilon \leq p^{v}(x,z^{v}) \leq \varepsilon^{-1}$ for all $v \in V$ and $x,z \in \mathbb{X}$. Let ν be a probability measure on $\mathbb{X}$, and suppose that*

$$\mathrm{Corr}(\nu,\beta) \leq \tfrac{1}{2} - (1-\varepsilon^2)e^{\beta(r+1)}\Delta$$

for a sufficiently small constant $\beta > 0$. Then

$$\sup_{x,z\in\mathbb{X}} |||(\mathsf{F}_s\nu)^{v}_{x,z} - (\tilde{\mathsf{F}}_s\nu)^{v}_{x,z}||| \leq 4e^{-\beta}(1-\varepsilon^{2\Delta})e^{-\beta d(v,\partial K(v))}$$

for every $s \in \mathbb{N}$ and every $K \in \mathcal{K}_{\sigma(s)}$ with $v \in K$. Furthermore,

$$|||\mathsf{F}_n\nu - \tilde{\mathsf{F}}_n\nu|||_J \leq 4|J|e^{-\beta}(1-\varepsilon^{2\Delta})e^{-\beta d(J,\partial K)}$$

for every $K \in \mathcal{K}_{\sigma(n)}$ and $J \subseteq K$.

Now we are ready to put all the results together and finalise the bound on the bias.

6.1 Proof of Theorem 2

Note $\tilde{\pi}_0^{\mu} = \mu$ and here we consider the case $\mu = \delta(x)$ for some $x \in \mathbb{X}$. Firstly, note

$$\varepsilon > \varepsilon_0 = \left(1 - 1/(18\Delta^2)\right)^{1/2\Delta} > \left(1 - 1/(16\Delta^2)\right)^{1/2\Delta}$$

and $\beta = -(2r)^{-1}\log 16\Delta^2(1-\varepsilon^{2\Delta}) > 0$ such that Lemma 6 holds. Moreover,

$$\mathrm{Corr}(\tilde{\pi}_n^{x},\beta) \;\leq\; 1/3 \;\leq\; \tfrac{1}{2} - (1-\varepsilon^2)e^{\beta(r+1)}\Delta \geq 4/9$$

for every $n \geq 0$ and Lemma 7 holds. Finally,

$$\mathrm{Corr}(\tilde{\pi}_n^{x},\beta) \;\leq\; 1/3 \leq 1/2 - 3(1-\varepsilon^{2\Delta})e^{2\beta r}\Delta^2 \geq 1/3$$

for every $n \geq 0$ and Lemma 1 holds. Thus, the bound on the decay of correlations holds and implies local filter stability and the one-step blocking approximation error bound all hold under the given parameter hypotheses.

Fix $K_{\sigma(s)} \in \mathcal{K}_{\sigma(s)}$ so $J \subseteq K_{\sigma(s)}$ for all $s \in \mathbb{N}$. Then

$$\begin{aligned}\|\!|\pi_n^x - \tilde{\pi}_n^x|\!\|_J &\leq \sum_{s=1}^{n} \|\!|\mathsf{F}_n \ldots \mathsf{F}_{s+1}\mathsf{F}_s\tilde{\pi}_{s-1}^x - \mathsf{F}_n \ldots \mathsf{F}_{s+1}\tilde{\mathsf{F}}_s\tilde{\pi}_{s-1}^x|\!\|_J \\ &\leq \|\!|\mathsf{F}_n\tilde{\pi}_{n-1}^x - \tilde{\mathsf{F}}_n\tilde{\pi}_{n-1}^x|\!\|_J + \sum_{s=1}^{n-1} \|\!|\mathsf{F}_n \ldots \mathsf{F}_{s+1}\mathsf{F}_s\tilde{\pi}_{s-1}^x - \mathsf{F}_n \ldots \mathsf{F}_{s+1}\tilde{\mathsf{F}}_s\tilde{\pi}_{s-1}^x|\!\|_J\end{aligned}$$

Following [14] by application of Lemma 1 and Lemma 6 we easily find

$$\|\!|\pi_n^x - \tilde{\pi}_n^x|\!\|_J \;\leq\; 8e^{-\beta}(1-\varepsilon^{2\Delta})\left(|J|e^{-\beta d(J,\partial K_{\sigma(n)})} + \sum_{s=1}^{n-1} e^{-\beta(n-s)} \sum_{v\in J} e^{-\beta d(v,\partial K_{\sigma(n)})}\right)$$

for every $n \geq 0$, $x \in \mathbb{X}$ and every $K_{\sigma(s)} \in \mathcal{K}_{\sigma(s)}$ such that $J \subseteq K_{\sigma(s)}$ for all $s \in \mathbb{N}$.

For any site v consider the sequence $K_{\sigma(s)}(v) \in \mathcal{K}_{\sigma(s)}$ for all $s \in \mathbb{N}$. The bound just given becomes

$$\|\!|\pi_n^x - \tilde{\pi}_n^x|\!\|_v \;\leq\; 8e^{-\beta}(1-\varepsilon^{2\Delta}) \sum_{s=1}^{n} e^{-\beta(n-s)} e^{-\beta d(v,\partial K_{\sigma(s)}(v))}$$

Recall the hypothesis $\sigma(s) = s \pmod m$, $s \in \mathbb{N}$. Now averaging the bound gives

$$\frac{1}{m}\sum_{k=0}^{m-1} \|\!|\pi_{n-k}^x - \tilde{\pi}_{n-k}^x|\!\|_v \;\leq\; 8e^{-\beta}(1-\varepsilon^{2\Delta})\frac{1}{m}\sum_{k=0}^{m-1}\sum_{s=1}^{n-k} e^{-\beta(n-k-s)} e^{-\beta d(v,\partial K_{\sigma(s)}(v))}$$

and we define $\|\!|\pi_{-s}^x - \tilde{\pi}_{-s}^x|\!\|_v$ for $s \in \mathbb{N}$ to be zero. Now

$$\begin{aligned}\frac{1}{m}\sum_{k=0}^{m-1}\sum_{s=1}^{n-k} e^{-\beta(n-k-s)} e^{-\beta d(v,\partial K_{\sigma(s)}(v))} &= \frac{1}{m}\sum_{k=0}^{m-1}\sum_{s=0}^{n-k-1} e^{-\beta s} e^{-\beta d(v,\partial K_{\sigma(n-k-s)}(v))} \\ &\leq \frac{1}{m}\sum_{k=0}^{m-1}\sum_{s=0}^{n-1} e^{-\beta s} e^{-\beta d(v,\partial K_{\sigma(n-k-s)}(v))}\end{aligned}$$

where now we extend $\sigma(s)$ to all $s \in \mathbb{Z}$ so $\sigma(s) = s \pmod m$. Then

$$\frac{1}{m}\sum_{k=0}^{m-1}\sum_{s=0}^{n-1} e^{-\beta s} e^{-\beta d(v,\partial K_{\sigma(n-k-s)}(v))} \;=\; \frac{1}{m}\sum_{s=0}^{n-1} e^{-\beta s}\sum_{k=0}^{m-1} e^{-\beta d(v,\partial K_{\sigma(n-k-s)}(v))} \;=\; \frac{1}{m}\sum_{s=0}^{n-1} e^{-\beta s}\sum_{k=0}^{m-1} e^{-\beta d(v,\partial K_{\sigma(k)}(v))}$$

where we swapped the summation order and used the fact that

$$\sum_{k=0}^{m-1} e^{-\beta d(v,\partial K_{\sigma(n-k-s)}(v))} = \sum_{k=0}^{m-1} e^{-\beta d(v,\partial K_{\sigma(k)}(v))}$$

for $s = 0$ and all $s \in \mathbb{N}$ because of the cyclical partitioning sequence. Now it follows that

$$\frac{1}{m}\sum_{k=0}^{m-1} \|\!|\pi_{n-k}^x - \tilde{\pi}_{n-k}^x|\!\|_v \;\leq\; \frac{8e^{-\beta}}{1-e^{-\beta}}(1-\varepsilon^{2\Delta})\frac{1}{m}\sum_{j=0}^{m-1} e^{-\beta d(v,\partial K_j(v))} \;=\; \frac{8e^{-\beta}}{1-e^{-\beta}}(1-\varepsilon^{2\Delta})\vartheta_m(v)$$

noting $\sum e^{-\beta s} \leq 1/(1-e^{-\beta})$. Let $\nabla_d(v) \triangleq \min_s d(v,\partial K_s)$ then the over bounding

$$\tfrac{8e^{-\beta}}{(1-e^{-\beta})}(1-\varepsilon^{2\Delta})\vartheta_m(v) \;\leq\; \tfrac{8e^{-\beta}}{(1-e^{-\beta})}(1-\varepsilon^{2\Delta})\exp\left[-\beta e^{-\beta(\Delta_d(v)-\nabla_d(v))}\frac{1}{m}\theta_m(v)\right]$$

follows from Slater's inequality and the proof is complete. □

Again, both inequalities in Theorem 2 imply that the bias introduced due to blocking can be spatially averaged (or smoothed) across a cyclical application of a sequence of partitions and both inequalities collapse to the result of [15] in the case $m = 1$.

References

[1] B.D.O. Anderson and J.B. Moore. *Optimal Filtering.* Prentice Hall, Englewood Cliffs, N.J., 1979.

[2] A. Beskos, D. Crisan, and A. Jasra. On the stability of sequential Monte Carlo methods in high dimensions. *The Annals of Applied Probability*, 24(4):1396–1445, August 2014.

[3] A. Beskos, D. Crisan, A. Jasra, and N. Whiteley. Error bounds and normalising constants for sequential Monte Carlo samplers in high dimensions. *Advances in Applied Probability*, 46(1):279–306, March 2014.

[4] P. Bickel, B. Li, and T. Bengtsson. Sharp failure rates for the bootstrap particle filter in high dimensions. In *Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh*, volume 3, pages 318–329. Institute of Mathematical Statistics, 2008.

[5] O. Cappé, E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*. Springer, New York, N.Y., 2005.

[6] D. Crisan and A. Doucet. A survey of convergence results on particle filtering methods for practitioners. *IEEE Transactions on Signal Processing*, 50(3):736–746, March 2002.

[7] P. Del Moral. *Mean Field Simulation for Monte Carlo Integration*. CRC Press, Boca Raton, FL, 2013.

[8] P. Del Moral and A. Guionnet. On the stability of interacting processes with applications to filtering and genetic algorithms. *Annales de l'Institut Henri Poincaré (B) Probability and Statistics*, 37(2):155–194, March 2001.

[9] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer, New York, N.Y., 2001.

[10] A.L. Gibbs and F.E. Su. On choosing and bounding probability metrics. *International Statistical Review*, 70(3):419–435, December 2002.

[11] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, 140(2):107–113, April 1993.

[12] Genshiro Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25, 1996.

[13] P.B. Quang, C. Musso, and F. Le Gland. An insight into the issue of dimensionality in particle filtering. In *Proceedings of the 13th International Conference on Information Fusion*, Edinburgh, Scotland, July 2010.

[14] P. Rebeschini and R. van Handel. Can local particle filters beat the curse of dimensionality? *arXiv e-print arXiv:1301.6585*, 2013.

[15] P. Rebeschini and R. van Handel. Comparison theorem for Gibbs measures. *arXiv e-print arXiv:1308.4117*, 2013.

[16] C. Snyder, T. Bengtsson, P. Bickel, and J. Anderson. Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136(12):4629–4640, December 2008.

[17] R. Van Handel. Uniform time average consistency of Monte Carlo particle filters. *Stochastic Processes and their Applications*, 119(11):3835–3861, November 2009.

[18] P.J. van Leeuwen. Particle filtering in geophysical systems. *Monthly Weather Review*, 137(12):4089–4114, December 2009.

[19] P.J. van Leeuwen. Nonlinear data assimilation in geosciences: An extremely efficient particle filter. *Quarterly Journal of the Royal Meteorological Society*, 136(653):1991–1999, October 2010.

[20] E. Vanmarcke. *Random Fields: Analysis and Synthesis*. World Scientific, Singapore, 2010.